

Pathwise Iteration for Backward SDEs

Christian Bender¹, Christian Gärtner¹, and Nikolaus Schweizer²

June 24, 2016

Abstract

We introduce a novel numerical approach for a class of stochastic dynamic programs which arise as discretizations of backward stochastic differential equations or semi-linear partial differential equations. Solving such dynamic programs numerically requires the approximation of nested conditional expectations, i.e., iterated integrals of previous approximations. Our approach allows us to compute and iteratively improve upper and lower bounds on the true solution starting from an arbitrary and possibly crude input approximation. We demonstrate the benefits of our approach in a high dimensional financial application.

Keywords: Backward stochastic differential equations, dynamic programming, iterated improvement, Monte Carlo

AMS subject classifications: 65C5, 65C30, 49L20, 93E20, 93E24

1 Introduction

Developing numerical methods for American option pricing, i.e. the optimal stopping problem, is one of the most specialized and developed fields in computational finance. We generalize several established numerical tools from optimal stopping to a class of convex stochastic dynamic programming equations. Applications include time discretization schemes for backward stochastic differential equations (BSDEs) or, equivalently [22], discretization schemes for semi-linear partial differential equations (PDEs) where the nonlinearity is convex (or concave). In these problems, the numerical challenge has its origin in a high order nesting of conditional expectation operators: The approximation at a given time step depends on iterated integrals over the approximations at all future time steps. The curse of dimensionality renders many numerical approaches infeasible in such a setting. This includes naive implementations of (nested) Monte Carlo. When moving from optimal stopping to BSDEs or semi-linear PDEs, a further numerical challenge arises as approximating derivatives becomes a necessity.

Our main contribution is a pathwise iteration approach which takes an approximate solution of the dynamic programming equation as an input and then constructs upper and lower confidence bounds on the true solution. Iteratively taking the super- and subsolutions corresponding to these bounds as inputs allows to refine the initial bounds. Thus, even a crude input approximation

¹Saarland University, Department of Mathematics, Postfach 151150, D-66041 Saarbrücken, Germany, bender@math.uni-sb.de; gaertner@math.uni-sb.de.

Financial support by the Deutsche Forschungsgemeinschaft under grant BE3933/5-1 is gratefully acknowledged.

²University of Duisburg-Essen, Mercator School of Management, Lotharstr. 65, D-46057 Duisburg, nikolaus.schweizer@uni-due.de.

– such as a constant function – may suffice to provably pin down the solution of a challenging high-dimensional problem up to a tight confidence interval.

For optimal stopping, such “primal-dual” approaches for the construction of upper and lower bounds go back to [21, 26, 25, 16, 1]. The approach was extended to our setting of convex dynamic programming equations in [4, 5], complementing the information relaxation approach of [7] which provides a generalization from optimal stopping to more general optimization problems. Iterative improvement methods of the primal-dual approach have been developed in [8] for the upper bounds and [18] for the lower bounds, building on earlier policy iteration techniques due to [17, 23]. We simplify and unify their arguments in terms of super- and subsolutions and generalize them beyond optimal stopping.

As a second contribution, we introduce two new methods, a minimization and a modified least-squares Monte Carlo (LSMC) method, for computing approximate solutions to stochastic dynamic programming equations. While we primarily use these methods as inputs for our improvement approach, both are of independent interest. In the minimization method, we use the pathwise recursion for the construction of upper bounds to minimize over a given family of generic input upper bounds. Methods of this type have been proposed in the stopping literature by [2, 9]. Our modified least-squares Monte Carlo algorithm builds on the regress-later method for optimal stopping [13] and its generalization to BSDEs, the martingale basis algorithm of [3]. Both approaches replace the true solution by an approximation as linear combination of basis functions, for which some computations can be performed in closed form. Our variant of the method has more modest requirements on what can be calculated explicitly (i.e. with negligible error), thus increasing its applicability and flexibility. Unlike in optimal stopping and as observed in [3], a considerable benefit of closed-form calculations is that they may allow to approximate derivatives without further error when derivatives of the basis functions are available. The goal of our modified LSMC algorithm is to gain flexibility by retaining only the availability of closed-form derivatives from these previous methods.

Both the minimization approach and the iterative improvement operate pathwise, i.e., trajectory by trajectory. Compared to classical LSMC methods [21, 20], they thus have a better scope for massively parallel implementations under memory constraints. See [15] for a recent contribution which highlights these issues and presents a variation of LSMC which is more amenable to parallelization. We confirm the practical applicability of our methods in a classical reference problem, pricing under funding risk in a financial market model driven by a five-dimensional Markov process. Depending on the time discretization, this corresponds to integrating out between 100 and 200 variables with a complex dependence structure in our Monte Carlo approach.

The paper is organized as follows: Section 2 introduces the setting. Section 3 develops the theory behind our iteration approach for subsolutions, i.e., lower bounds, while Section 4 provides the analogous results for upper bounds. Section 5 provides an overview of our numerical approach, including our new approximation methods. Numerical results in the context of funding risk are presented in Section 6.

2 Setup

Throughout the paper, we study the following type of convex dynamic programming equation on a complete filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_j)_{j=0, \dots, J}, P)$ in discrete time:

$$Y_j = F_j(E_j[\beta_{j+1}Y_{j+1}]), \quad j = J-1, \dots, 0, \quad Y_J = \xi, \quad (1)$$

with given data ξ , F and β (to be specified below), where $E_j[\cdot]$ denotes the conditional expectation with respect to \mathcal{F}_j . This type of recursive equation encompasses the dynamic programming equation for optimal stopping of an adapted discrete time process S (or, in financial terms, the Bermudan option pricing problem),

$$Y_j = \max\{S_j, E_j[Y_{j+1}]\}, \quad Y_J = S_J, \quad (2)$$

see e.g. [18], and discretization schemes for backward stochastic differential equations of the form

$$Y_j = E_j[Y_{j+1}] + (t_{j+1} - t_j) G\left(t_j, E_j[Y_{j+1}], E_j\left[\frac{W_{t_{j+1}} - W_{t_j}}{t_{j+1} - t_j} Y_{j+1}\right]\right), \quad Y_J = \xi, \quad (3)$$

for given data ξ and G . Here, (t_0, \dots, t_J) denotes a partition of a time interval $[0, T]$, W is a multidimensional Brownian motion (whose increments may be truncated for practical purposes), and \mathcal{F}_j is the information generated by W up to time t_j , see [11] for this specific scheme in the more general context of second order BSDEs and [3] for a literature overview.

We assume that $F_j : \Omega \times \mathbb{R}^D \rightarrow \mathbb{R}$ is measurable for every $j = 0, \dots, J-1$ and that the process $(j, \omega) \mapsto F_j(\omega, z)$ is adapted for every $z \in \mathbb{R}^D$. Moreover, for every $j = 0, \dots, J-1$ and $\omega \in \Omega$, the map $z \mapsto F_j(\omega, z)$ is convex in z . Additionally, F_j satisfies a (stochastic) polynomial growth condition for every $j = 0, \dots, J-1$, i.e. there exist a constant $q \geq 0$ and an adapted, nonnegative processes α , which is in $L^p(\Omega, P)$ for every $p \geq 1$, such that

$$|F_j(z)| \leq \alpha_j(1 + |z|^q)$$

holds P -a.s. for every $z \in \mathbb{R}^D$. The \mathbb{R}^D -valued process β is adapted and in $L^p(\Omega, P)$ for every $p \geq 1$. The terminal condition ξ is an \mathcal{F}_J -measurable, \mathbb{R} -valued random variable with $E[|\xi|^p] < \infty$ for all $p \geq 1$.

Further, we introduce the following notation: For $m \in \mathbb{N}$, we denote by $L^{\infty-}(\mathbb{R}^m)$ the set of \mathbb{R}^m -valued random variables that are in $L^p(\Omega, P)$ for all $p \geq 1$. The set of \mathcal{F}_j -measurable random variables that are in $L^{\infty-}(\mathbb{R}^m)$ is denoted by $L_j^{\infty-}(\mathbb{R}^m)$. In addition, $L_{ad}^{\infty-}(\mathbb{R}^m)$ denotes the set of adapted processes Z such that $Z_j \in L_j^{\infty-}(\mathbb{R}^m)$ for every $j = 0, \dots, J$. From the integrability properties of the terminal condition ξ and the weight process β as well as the polynomial growth condition on F , we deduce by backward induction that the (P -a.s. unique) solution Y to (1) is in $L_{ad}^{\infty-}(\mathbb{R})$.

Super- and subsolutions which are central later on are defined as follows:

Definition 2.1. A process Y^{up} (respectively Y^{low}) $\in L_{ad}^{\infty-}(\mathbb{R})$ is called *supersolution* (respectively *subsolution*) to the dynamic program (1) if $Y_j^{up} \geq Y_j$ (respectively $Y_j^{low} \leq Y_j$) and for every $j = 0, \dots, J-1$ it holds that

$$Y_j^{up} \geq F_j\left(E_j\left[\beta_{j+1} Y_{j+1}^{up}\right]\right), \quad P\text{-a.s.}$$

(and with ' \geq ' replaced by ' \leq ' for a subsolution).

In general, we cannot expect that super- and subsolutions Y^{up} and Y^{low} to (1) are bounds on the true solution Y , i.e. we need not have that $Y_j^{up} \geq Y_j \geq Y_j^{low}$ holds P -a.s. for every $j = 0, \dots, J$. To ensure this, we impose the following *monotonicity* assumption throughout this paper: For $Y^{(1)}, Y^{(2)} \in L^{\infty-}(\mathbb{R})$ with $Y^{(1)} \geq Y^{(2)}$ P -a.s. and every $j = 0, \dots, J-1$, it holds that

$$F_j(\beta_{j+1} Y^{(1)}) \geq F_j(\beta_{j+1} Y^{(2)}), \quad P\text{-a.s.} \quad (4)$$

Applying Theorem 4.3 in [5] twice (first with the filtration $(\mathcal{G}_j)_{j=0,\dots,J} = (\mathcal{F})_{j=0,\dots,J}$, and then with the given filtration $(\mathcal{F}_j)_{j=0,\dots,J}$), we observe that this monotonicity assumption implies the following *comparison principle*:

Proposition 2.2. *Let Y^{up} and Y^{low} be super- and subsolutions to (1). Then, under the given assumptions, it holds that, for every $j = 0, \dots, J$,*

$$Y_j^{up} \geq Y_j^{low} \quad P\text{-a.s.}$$

The improvement algorithms presented in the following sections are based on the primal-dual methodology introduced by [1, 16, 25] in the context of Bermudan option pricing and further developed in [4] and [5] for dynamic programming equations of the form (1). This approach relies on the choice of suitable martingales and controls which are derived from an approximate solution to (1) and are used as an input for constructing super- and subsolutions. We, therefore, denote by \mathcal{M}_D the set of \mathbb{R}^D -valued martingales which are elements of $L_{ad}^{\infty-}(\mathbb{R}^D)$. For a process $Z \in L_{ad}^{\infty-}(\mathbb{R}^D)$, we refer to the martingale part of the Doob decomposition of Z , which is given by

$$\sum_{i=0}^{j-1} Z_{i+1} - E_i[Z_{i+1}], \quad j = 0, \dots, J,$$

as *Doob martingale* of Z . In particular, we get that the Doob martingale of the process $\beta \bar{Z}$ is in \mathcal{M}_D for any $\bar{Z} \in L_{ad}^{\infty-}(\mathbb{R})$. While suitable martingales are the main ingredient of the upper bounds, we derive lower bounds by rewriting (1) as a stochastic control problem using convex duality techniques. To this end, recall that the convex conjugate of F_j is, for every $\omega \in \Omega$, given by

$$F_j^\#(\omega, u) := \sup_{z \in \mathbb{R}^D} (u^\top z - F_j(\omega, z)),$$

with effective domain

$$D_{F^\#}^{(j,\omega)} = \left\{ u \in \mathbb{R}^D \mid F_j^\#(\omega, u) < \infty \right\}.$$

As we will see below, the sets of admissible controls in our problem are given by

$$\mathcal{A}_j = \left\{ (r_i)_{i=j,\dots,J-1} \mid r_i \in L_i^{\infty-}(\mathbb{R}^D), F_i^\#(r_i) \in L^{\infty-}(\mathbb{R}) \text{ for } i = j, \dots, J-1 \right\}, \quad (5)$$

where $j = 0, \dots, J-1$.

3 Improvement of subsolutions

In this section we propose an iterative algorithm to improve a given subsolution to (1). This approach generalizes in some sense the idea of [18], who presented an iterative method to improve a given family of stopping times in the context of Bermudan option pricing. We begin this section by recalling a construction of subsolutions from [4] and, then, explain how it can be used to improve arbitrary subsolutions.

In order to construct a subsolution to (1), we linearize this dynamic programming equation in the following way: By convexity and closedness of F_j , we have due to Theorem 12.2 in [24]

that $F_j^{\#\#} = F_j$ for every $j = 0, \dots, J-1$ and $\omega \in \Omega$. Hence, for every $j = 0, \dots, J-1$, $\omega \in \Omega$ and $z \in \mathbb{R}^D$, it holds that

$$F_j(\omega, z) = \sup_{u \in \mathbb{R}^D} u^\top z - F_j^\#(u). \quad (6)$$

From Lemma A.1 in [5], we get existence of an adapted process $r^* \in \mathcal{A}_0$ which solves

$$(r_i^*)^\top E_i[\beta_{i+1} Y_{i+1}] - F_i^\#(r_i^*) = F_i(E_i[\beta_{i+1} Y_{i+1}]) \quad P\text{-a.s.}, \quad (7)$$

for every $i = 0, \dots, J-1$. We now define the typically non-adapted process θ^{low} as in Remark 3.6 (ii) in [4]. To this end, we fix a martingale $M \in \mathcal{M}_D$ and an admissible control $r \in \mathcal{A}_0$. Then, the pathwise recursion for $\theta^{low} := \theta^{low}(r, M)$ is, in our notation, given as follows:

$$\theta_j^{low} = r_j^\top \beta_{j+1} \theta_{j+1}^{low} - r_j^\top \Delta M_{j+1} - F_j^\#(r_j), \quad j = J-1, \dots, 0, \quad \theta_J^{low} = \xi, \quad (8)$$

where $\Delta M_{j+1} := M_{j+1} - M_j$. By backward induction, we get that $\theta_j^{low} \in L^{\infty-}(\mathbb{R})$ for every $j = 0, \dots, J$, since r and β are in $L_{ad}^{\infty-}(\mathbb{R}^D)$ and $M \in \mathcal{M}_D$ by assumption. Hence, we can define a subsolution Y^{low} by $Y_j^{low} := E_j[\theta_j^{low}]$ for every $j = 0, \dots, J$. Indeed, by the tower property of the conditional expectation and (6), we observe that

$$Y_j^{low} = r_j^\top E_j[\beta_{j+1} \theta_{j+1}^{low}] - F_j^\#(r_j) \leq F_j\left(E_j\left[\beta_{j+1} Y_{j+1}^{low}\right]\right)$$

holds and, hence, by Proposition 2.2, we conclude that $Y_j \geq Y_j^{low}$ P -a.s., for any $j = 0, \dots, J$. Moreover, [4] prove that the solution Y to (1) is the value of a primal maximization problem, i.e.

$$Y_j = \text{esssup}_{r \in \mathcal{A}_j} E_j[\theta_j^{low}(r, M)] \quad P\text{-a.s.}, \quad j = 0, \dots, J.$$

Indeed, every control $r^* \in \mathcal{A}_j$ which satisfies (7) for $i = j, \dots, J-1$ achieves the maximum. We emphasize that the expression $E_j[\theta_j^{low}(r, M)]$ is independent of M , as the martingale is only a control variate in the recursion for θ^{low} and thus vanishes by taking conditional expectation. However, a straightforward computation shows that the Doob martingale M^* of βY acts as a perfect control variate in the case of optimal controls, i.e. for every $j = 0, \dots, J$ it holds that,

$$\theta_j^{low}(r^*, M^*) = Y_j \quad P\text{-a.s.} \quad (9)$$

Iterative improvement of subsolutions

Suppose we are given an arbitrary subsolution \bar{Y} . We next show that the construction of the process $\theta^{low}(r, M)$ in (8), implies an improvement of the subsolution \bar{Y} in the sense that for suitable choices of $r \in \mathcal{A}_0$ and $M \in \mathcal{M}_D$ we have

$$Y_j \geq E_j\left[\theta_j^{low}(r, M)\right] \geq \bar{Y}_j \quad P\text{-a.s.}$$

for every $j = 0, \dots, J$. The subsolution $(E_j[\theta_j^{low}(r, M)])_{j=0, \dots, J}$ is then called an *improvement* of the subsolution \bar{Y} . Theorem 3.1 below, explains how to construct such an improvement. Further, we show that our construction only gets stuck if the subsolution \bar{Y} , which we want to improve, already coincides with the solution Y to (1).

Theorem 3.1. Let $j \in \{0, \dots, J-1\}$, let \bar{Y} be a subsolution to (1) and denote by $\bar{M} \in \mathcal{M}_D$ the Doob martingale of $\beta\bar{Y}$. Further let $\bar{r} \in \mathcal{A}_0$ be an adapted process that solves

$$\bar{r}_i^\top E_i [\beta_{i+1} \bar{Y}_{i+1}] - F_i^\#(\bar{r}_i) = F_i(E_i [\beta_{i+1} \bar{Y}_{i+1}]) \quad P\text{-a.s.} \quad (10)$$

for every $i = 0, \dots, J-1$. Then, for any $M \in \mathcal{M}_D$, $\theta^{low}(\bar{r}, M)$ defined by (8) satisfies

$$Y_i \geq E_i [\theta_i^{low}(\bar{r}, M)] \geq F_i(E_i [\beta_{i+1} \bar{Y}_{i+1}]) \geq \bar{Y}_i \quad P\text{-a.s.}, \quad (11)$$

for all $i = 0, \dots, J-1$. Moreover, if $\bar{Y}_i = Y_i$ for all $i = j+1, \dots, J$, then

$$E_j [\theta_j^{low}(\bar{r}, M)] = \theta_j^{low}(\bar{r}, \bar{M}) = Y_j \quad P\text{-a.s.}$$

Proof. As we have seen above, the process $(E_j[\theta_j^{low}(r, M)])_{j=0, \dots, J}$ defines a subsolution for any martingale $M \in \mathcal{M}_D$ and $r \in \mathcal{A}_0$, so that the first inequality in (11) is already shown. The last inequality in (11) is immediate, as \bar{Y} is assumed to be a subsolution. To prove the remaining inequality in (11), we denote $\theta^{low} = \theta^{low}(\bar{r}, \bar{M})$. Recalling that $E_j[\theta_j^{low}(r, M)]$ does not depend on the choice of $M \in \mathcal{M}_D$, it suffices to show that

$$\theta_i^{low} \geq F_i(E_i [\beta_{i+1} \bar{Y}_{i+1}]) \quad P\text{-a.s.}$$

by backward induction on i . The assertion then follows by the monotonicity of the conditional expectation. The case $i = J$ is trivial, since we have $\theta_J^{low} = \xi \geq \bar{Y}_J$, by definition. Now suppose that the assertion is true for $i+1 \in \{1, \dots, J\}$, and, thus, we have $\theta_{i+1}^{low} \geq \bar{Y}_{i+1}$ P -a.s. Then, it follows from the definition of \bar{M} , (10), and the induction hypothesis that

$$\begin{aligned} \theta_i^{low} &= \bar{r}_i^\top (\beta_{i+1} \theta_{i+1}^{low} - (\beta_{i+1} \bar{Y}_{i+1} - E_i[\beta_{i+1} \bar{Y}_{i+1}])) - F_i^\#(\bar{r}_i) \\ &= \bar{r}_i^\top \beta_{i+1} (\theta_{i+1}^{low} - \bar{Y}_{i+1}) + F_i(E_i[\beta_{i+1} \bar{Y}_{i+1}]) \geq F_i(E_i[\beta_{i+1} \bar{Y}_{i+1}]). \end{aligned}$$

Here, the inequality is a consequence of the induction hypothesis and the positivity of $\bar{r}_i^\top \beta_{i+1}$ which is due to the monotonicity assumption (4) on the function F_i , see Theorem 4.3 in [5] (with the constant full information filtration $(\mathcal{G}_j)_{j=0, \dots, J} = (\mathcal{F})_{j=0, \dots, J}$).

To complete the proof, we assume that $\bar{Y}_i = Y_i$ for all $i = j+1, \dots, J$. Hence, we observe by (10), that \bar{r}_i satisfies the optimality condition (7) P -a.s. for every $i = j, \dots, J$. Further, we conclude by the definition of \bar{M} that the increments $\bar{M}_{i+1} - \bar{M}_i$ coincide with the ones of the Doob martingale of βY for $i = j, \dots, J-1$. By (9), we thus obtain that $\theta_j^{low} = Y_j$ P -a.s., which completes the proof. \square

When starting with an arbitrary subsolution, we typically do not obtain the solution Y by applying the approach described in Theorem 3.1 once. However, this construction can be iterated in a straightforward way: Let $Y^{(low, 0)}$ be a subsolution and define $\theta^{(low, 0)} := Y^{(low, 0)}$. We define the k -th iteration according to (8) by

$$\theta^{(low, k)} := \theta^{low}(r^{(k)}, M^{(k)}), \quad k \geq 1, \quad (12)$$

where the process $r^{(k)} \in \mathcal{A}_0$ is for every $j = 0, \dots, J$ given by

$$\left(r_j^{(k)}\right)^\top E_j [\beta_{j+1} \theta_{j+1}^{(low, k-1)}] - F_j^\#(r_j^{(k)}) = F_j(E_j [\beta_{j+1} \theta_{j+1}^{(low, k-1)}]), \quad (13)$$

and $M^{(k)} \in \mathcal{M}_D$ is arbitrary. Applying Theorem 3.1 iteratively, we observe that $E_j[\theta_j^{(low,k)}] \geq E_j[\theta_j^{(low,k-1)}]$, P -a.s., for every $k \geq 1$ and $j = 0, \dots, J$. Moreover,

$$E_i \left[\theta_i^{(low, J-j)} \right] = Y_i \quad P\text{-a.s.},$$

whenever $i \geq j$. In the last equation, the conditional expectation on the left-hand side can be removed, when each $M^{(k)}$ is taken as the Doob martingale of $\beta_j E_j[\theta_j^{(low, k-1)}]$. We, thus, observe that Y is the P -a.s. unique fixed point of this iteration, which actually terminates after at most J iteration steps.

Improvement of a family of subsolutions

In Section 5 below, we explain that the numerical costs of algorithms based on (12) tend to grow exponentially in the number of iterations k . For this reason, a moderate number of iterations must suffice in practical implementations. One way to address this issue is to improve a whole family of subsolutions simultaneously instead of just one subsolution. To this end, let $(\bar{Y}^{\{l\}})_{l \in I}$ be a family of subsolutions, where I is a finite index set. Further, we denote by $K(j)$, $j = 1, \dots, J$, a nondecreasing sequence of subsets of I , i.e. it holds that $K(j) \subset K(j+1)$. Then, we consider the predictable, I -valued process

$$l^*(j) = \inf \left\{ l \in K(j) \mid \forall l \in K(j) \ F_{j-1} \left(E_{j-1} \left[\beta_j \bar{Y}_j^{\{l\}} \right] \right) \geq F_{j-1} \left(E_{j-1} \left[\beta_j \bar{Y}_j^{\{l\}} \right] \right) \right\}.$$

This means that, at every time point $j = 1, \dots, J$, we only consider those subsolutions which are represented in the subset $K(j)$ and the random variable $l^*(j)$ returns an index $l \in K(j)$ at which the evaluation of F_{j-1} is maximized. In the simplest case $K(j) = I$ for all $j = 1, \dots, J$. More sophisticated choices of $K(j)$ allow to reduce the computational costs to determine l^* . We claim that the process \bar{Y} which is given by

$$\bar{Y}_j = \bar{Y}_j^{\{l^*(j)\}} \mathbf{1}_{\{j>0\}} + F_0(E_0[\beta_1 \bar{Y}_1]) \mathbf{1}_{\{j=0\}} \quad (14)$$

is a subsolution to (1), which allows us to improve the subsolutions $(\bar{Y}^{\{l\}})_{l \in I}$ simultaneously. To examine the subsolution property of \bar{Y} , we first observe that the case $j = 0$ is trivial, since we have $\bar{Y}_0 = F_0(E_0[\beta_1 \bar{Y}_1])$ by definition. For the case $j > 0$, we get by the subsolution property of $\bar{Y}^{\{l\}}$ for every $l \in I$, and as $K(j) \subset K(j+1)$, that

$$\begin{aligned} \bar{Y}_j &= \sum_{l \in K(j)} \bar{Y}_j^{\{l\}} \mathbf{1}_{\{l^*(j)=l\}} \leq \sum_{l \in K(j)} F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1}^{\{l\}} \right] \right) \mathbf{1}_{\{l^*(j)=l\}} \\ &\leq \sum_{l \in K(j)} F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1}^{\{l^*(j+1)\}} \right] \right) \mathbf{1}_{\{l^*(j)=l\}} = F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1} \right] \right) \quad P\text{-a.s.} \end{aligned}$$

Hence, Theorem 3.1 can be applied to the process \bar{Y} and implies, for $\theta^{low} = \theta^{low}(\bar{r}, M)$,

$$E_j \left[\theta_j^{low} \right] \geq F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1} \right] \right) = \max_{l \in K(j+1)} F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1}^{\{l\}} \right] \right) \geq \max_{l \in K(j+1)} \bar{Y}_j^{\{l\}} \quad (15)$$

P -a.s. for all $j = 0, \dots, J-1$, where \bar{r} is for every $j = 0, \dots, J-1$ given by (10) and where $M \in \mathcal{M}_D$. Thus, if $K(j) = I$ for all $j = 1, \dots, J$, we achieve a simultaneous improvement of all subsolutions $(\bar{Y}^{\{l\}})_{l \in I}$ by improving \bar{Y} .

Example 3.2. We consider the optimal stopping problem of an adapted process S , whose value process Y (the so-called Snell envelope) is governed by (2). Suppose $(\tau_l)_{l \in I}$, where $I = \{0, \dots, J\}$, is a family of I -valued stopping times which is consistent in the sense of [18]: For every $l \in I$

$$\tau_l \geq l \quad \text{and} \quad (\tau_l > l \Rightarrow \tau_l = \tau_{l+1}).$$

This consistency condition implies that each of the processes $\bar{Y}_j^{\{l\}} := E_j[S_{\max\{\tau_l, \tau_j\}}]$, $l \in I$, defines a subsolution to (2). Define $K(j) = \{0, \dots, \min\{j + \kappa - 1, J\}\}$ for some window parameter $\kappa \in \mathbb{N}$. Specializing the improvement condition (10) based on the subsolution constructed in (14) to the optimal stopping problem, one can verify that the improved subsolution satisfies, for every control variate $M \in \mathcal{M}_D$, $E_j[\theta^{low}(\bar{r}, M)] = E_j[S_{\bar{\tau}_j}]$ where the stopping times $\bar{\tau}_j$ are given by

$$\bar{\tau}_j = \inf\{i \geq j; S_i \geq \max_{l=i+1, \dots, \min\{i+\kappa, J\}} E_i[S_{\tau_l}]\}, \quad j = 0, \dots, J.$$

Hence, for every $j = 0, \dots, J$, by (15) and the consistency condition,

$$E_j[S_{\bar{\tau}_j}] \geq \max_{l=j+1, \dots, \min\{j+\kappa, J\}} \max\{E_j[S_{\tau_l}], S_j\} \geq \max_{l=j, \dots, \min\{j+\kappa, J\}} E_j[S_{\tau_l}], \quad P\text{-a.s.}$$

Thus, we recover the policy improvement result in Theorem 3.1 of [18] as a special case of our approach.

4 Improvement of supersolutions

In this section, we propose an iterative way for improving supersolutions to convex dynamic programs like (1). We generalize the construction of [8], who presented an improvement approach for supersolutions in the context of optimal stopping. Similar to Section 3, we build our approach on the pathwise recursion for upper bounds presented in [4]. Therefore, we begin this section with a brief overview of their construction and explain how it can be applied for improving arbitrary supersolutions. The remainder of this section is dedicated to transferring the results obtained in Section 3 to supersolutions.

The main idea of the pathwise approach presented in [4] is to remove the appearing conditional expectations in (1) and, instead, subtract a martingale increment. More precisely, let $M \in \mathcal{M}_D$ be a martingale. Then, we define the typically non-adapted process $\theta^{up} := \theta^{up}(M)$ by

$$\theta_j^{up} = F_j(\beta_{j+1}\theta_{j+1}^{up} - \Delta M_{j+1}), \quad j = 0, \dots, J-1, \quad \theta_J^{up} = \xi. \quad (16)$$

Due to the polynomial growth condition on F_j and the integrability properties of β and M , we get by backward induction that $\theta_j^{up} \in L^\infty(\mathbb{R})$ for every $j = 0, \dots, J$. Setting $(Y_j^{up})_{j=0, \dots, J} = (E_j[\theta_j^{up}])_{j=0, \dots, J}$, we observe immediately by Jensen's inequality, the martingale property of M , and the tower property of the conditional expectation that Y^{up} is a supersolution to (1):

$$Y_j^{up} = E_j[F_j(\beta_{j+1}\theta_{j+1}^{up} - \Delta M_{j+1})] \geq F_j(E_j[\beta_{j+1}Y_{j+1}^{up}]).$$

Similarly to the results in Section 3, [4] show that the solution Y to (1) can be represented as a dual minimization problem, i.e.

$$Y_j = \operatorname{ess\,inf}_{M \in \mathcal{M}_D} E_j[\theta_j^{up}(M)] \quad P\text{-a.s.} \quad (17)$$

for every $j = 0, \dots, J$. The Doob martingale M^* of βY achieves the minimum and additionally is even pathwise optimal, i.e. for every $j = 0, \dots, J$,

$$Y_j = \theta_j^{up}(M^*) \quad P\text{-a.s.} \quad (18)$$

Iterative improvement of supersolutions

As in Section 3, we show that an arbitrary supersolution \bar{Y} can be improved in the sense that for a suitable martingale $M \in \mathcal{M}_D$ and the process $\theta^{up}(M)$ from (16) the inequality

$$Y_j \leq E_j[\theta_j^{up}(M)] \leq \bar{Y}_j$$

holds P -a.s. for every $j = 0, \dots, J$. In the context of optimal stopping, [8] show that taking the Doob martingale of a given supersolution, leads to an improvement. Theorem 4.1 generalizes this idea to convex dynamic programs of the form (1):

Theorem 4.1. *Let $j \in \{0, \dots, J-1\}$ and let \bar{Y} be a supersolution to (1). Further, let \bar{M} be the Doob martingale of the process $\beta\bar{Y}$. Then, the process $\theta^{up}(\bar{M})$ satisfies*

$$Y_i \leq E_i[\theta_i^{up}(\bar{M})] \leq F_i(E_i[\beta_{i+1}\bar{Y}_{i+1}]) \leq \bar{Y}_i \quad P\text{-a.s.} \quad (19)$$

for all $i = 0, \dots, J$. Moreover, if $\bar{Y}_i = Y_i$ for all $i = j+1, \dots, J$, then

$$\theta_j^{up}(\bar{M}) = Y_j \quad P\text{-a.s.}$$

Proof. The overall strategy of proof is similar to the one of Theorem 3.1. At the beginning of this section, we have already shown that $(E_j[\theta_j^{up}(M)])_{j=0, \dots, J}$ is a supersolution for any martingale $M \in \mathcal{M}_D$, which yields the first inequality in (19). The last one is due to the supersolution property of \bar{Y} . To show the remaining inequality, we prove again the slightly stronger assertion

$$\theta_i^{up} := \theta_i^{up}(\bar{M}) \leq F_i(E_i[\beta_{i+1}\bar{Y}_{i+1}]) \quad P\text{-a.s.}, \quad i = 0, \dots, J-1,$$

from which we obtain (19) by the monotonicity of the conditional expectation. The proof is by backward induction on i , with the case $i = J$ being trivial, since, by definition, we have $\theta_J^{up} = \xi \leq \bar{Y}_J$. Now suppose that the assertion is true for $i+1 \in \{1, \dots, J\}$, i.e. $\theta_{i+1}^{up} \leq \bar{Y}_{i+1}$ P -a.s. Hence, we conclude by the definition of \bar{M} , the monotonicity assumption (4), and the induction hypothesis that

$$\begin{aligned} \theta_i^{up} &= F_i(\beta_{i+1}\theta_{i+1}^{up} - (\beta_{i+1}\bar{Y}_{i+1} - E_i[\beta_{i+1}\bar{Y}_{i+1}])) \\ &\leq F_i(\beta_{i+1}\bar{Y}_{i+1} - (\beta_{i+1}\bar{Y}_{i+1} - E_i[\beta_{i+1}\bar{Y}_{i+1}])) = F_i(E_i[\beta_{i+1}\bar{Y}_{i+1}]). \end{aligned}$$

Here, we exploit that $z \mapsto F_i(z - (\beta_{i+1}\bar{Y}_{i+1} - E_i[\beta_{i+1}\bar{Y}_{i+1}]))$ inherits the monotonicity property (4) from F_i by the equivalent characterization of the monotonicity property via positivity in Theorem 4.3 of [5] (with the constant full information filtration), because its convex conjugate is given by $F_i^\#(u) + u^\top(\beta_{i+1}\bar{Y}_{i+1} - E_i[\beta_{i+1}\bar{Y}_{i+1}])$.

To complete the proof, we assume that $\bar{Y}_i = Y_i$ for all $i = j+1, \dots, J$, where $j \in \{0, \dots, J-1\}$ is fixed. Then, again, the increments $\bar{M}_{i+1} - \bar{M}_i$ coincide with those of the Doob martingale M^* of βY for $i = j, \dots, J-1$. Hence, (18) concludes. \square

As in Section 3, this improvement can be iterated several times. For a given supersolution $Y^{(up,0)}$ define $\theta^{(up,0)} := Y^{(up,0)}$ and define $\theta^{(up,k)}$ according to (16) by

$$\theta^{(up,k)} := \theta^{up}(M^{(k)}), \quad k \geq 1, \quad (20)$$

where each $M^{(k)}$ is given by

$$M_j^{(k)} = \sum_{i=0}^{j-1} \beta_{i+1} E_{i+1} \left[\theta_{i+1}^{(up,k-1)} \right] - E_i \left[\beta_{i+1} \theta_{i+1}^{(up,k-1)} \right], \quad j = 0, \dots, J. \quad (21)$$

Then, iterative application of Theorem 4.1 yields $E_j[\theta_j^{(up,k)}] \leq E_j[\theta_j^{(up,k-1)}]$, P -a.s., for every $k \geq 1$ and $j = 0, \dots, J$, and

$$\theta_i^{(up,J-j)} = Y_i, \quad P\text{-a.s.},$$

whenever $i \geq j$. So the upper bound iteration also terminates after at most J steps at the true solution Y .

Improvement of a family of supersolutions

At the end of Section 3 we explained how to improve a given family of subsolutions. The same idea can be applied here in order to simultaneously improve a family of supersolutions $(\bar{Y}^{\{l\}})_{l \in I}$, where I is a finite index set. We now consider the predictable, I -valued process

$$l_*(j) = \inf \left\{ l \in K(j) \mid \forall \kappa \in K(j) \ F_{j-1} \left(E_{j-1} \left[\beta_j \bar{Y}_j^{\{l\}} \right] \right) \leq F_{j-1} \left(E_{j-1} \left[\beta_j \bar{Y}_j^{\{\kappa\}} \right] \right) \right\}$$

for every $j = 1, \dots, J$, where $K(j)$ is again a nondecreasing family of subsets of I . Then, the process \bar{Y} defined by

$$\bar{Y}_j = \bar{Y}_j^{\{l_*(j)\}} \mathbf{1}_{\{j > 0\}} + F_0 \left(E_0 \left[\beta_1 \bar{Y}_1 \right] \right) \mathbf{1}_{\{j=0\}}$$

is, by similar arguments as in Section 3, a supersolution to (1). Thus, by Theorem 4.1,

$$E_j \left[\theta_j^{up}(\bar{M}) \right] \leq F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1} \right] \right) = \min_{l \in K(j+1)} F_j \left(E_j \left[\beta_{j+1} \bar{Y}_{j+1}^{\{l\}} \right] \right) \leq \min_{l \in K(j+1)} \bar{Y}_j^{\{l\}}$$

P -a.s. for every $j = 0, \dots, J-1$, where \bar{M} denotes the Doob martingale of $\beta \bar{Y}$. Hence, in the case $K(j) = I$ for $j = 1, \dots, J$, improving \bar{Y} results again in a simultaneous improvement of all supersolutions $(\bar{Y}^{\{l\}})_{l \in I}$.

5 Implementation

In this section, we explain how to implement algorithms based on the iterative improvement approaches of Sections 3 and 4. In order to transform these results into implementable algorithms, one needs to construct a sub- and a supersolution as input. Moreover, the conditional expectations which appear in the iterative constructions of the controls in (13) and the Doob martingales in (21) must be approximated numerically. For the numerical approximation of the conditional expectation within the iterative improvement we apply, as in [18], a plain Monte Carlo implementation. In contrast to a naive plain Monte Carlo implementation of the dynamic programming equation (1) (which leads to infeasible J nested layers of simulation), the number of layers of simulation in the iterative improvement algorithm depends on the number of iteration steps which are performed. As we shall demonstrate in the numerical examples, two improvement steps are feasible, when the input super- and subsolutions are available in closed form. Therefore, we focus on the construction of closed-form inputs in Section 5.1, before we explain the somewhat standard nested simulation approach for the iterative improvement in Section 5.2.

As a first step, however, we specialize to the following Markovian framework: We assume that $(B_j)_{j=0,\dots,J}$ is an $\mathbb{R}^{\mathcal{D}}$ -dimensional adapted process (with $\mathcal{D} \geq D$), such that the first D components of B_j are given by β_j and B_j is independent of \mathcal{F}_{j-1} , for every $j = 1, \dots, J$. X is supposed to be an \mathbb{R}^N -valued Markovian process of the form

$$X_j = h_j(X_{j-1}, B_j), \quad j = 1, \dots, J, \quad (22)$$

for measurable functions $h_j : \mathbb{R}^N \times \mathbb{R}^{\mathcal{D}} \rightarrow \mathbb{R}^N$, starting at $X_0 = x_0 \in \mathbb{R}^N$. This forward equation for the state process X could arise, e.g., as a time discretization of a stochastic differential equation. Moreover, for the generator F_j of the dynamic program (1) we assume existence of measurable functions $f_j : \mathbb{R}^N \times \mathbb{R}^{\mathcal{D}} \rightarrow \mathbb{R}$ satisfying $F_j(\cdot) = f_j(X_j, \cdot)$, i.e., F_j depends on ω only through the Markovian process X . Then, we consider a Markovian version of the dynamic program (1) in the form

$$Y_j = f_j(X_j, E_j[\beta_{j+1} Y_{j+1}]), \quad j = 0, \dots, J-1, \quad Y_J = g(X_J), \quad (23)$$

where $g : \mathbb{R}^N \rightarrow \mathbb{R}$ is measurable. In this framework, Y_j is a deterministic function of X_j (and, in particular, Y_0 is a constant). In view of (22), we obtain, for every $j = 1, \dots, J$, a measurable function $y_j : \mathbb{R}^N \times \mathbb{R}^{\mathcal{D}} \rightarrow \mathbb{R}$ such that $Y_j = y_j(X_{j-1}, B_j)$. Denoting by P_{B_j} the law of B_j , we can, thus, write $E_j[\beta_{j+1} Y_{j+1}] = z_j(X_j)$ with

$$z_j(x) = \left(\int_{\mathbb{R}^{\mathcal{D}}} b_1 y_{j+1}(x, b) P_{B_{j+1}}(db), \dots, \int_{\mathbb{R}^{\mathcal{D}}} b_D y_{j+1}(x, b) P_{B_{j+1}}(db) \right)^\top.$$

5.1 Computation of the input sub- and supersolution

For the construction of the input sub- and supersolutions, we first approximate y_j by a linear combination of a given set of basis functions $\eta_j^1, \dots, \eta_j^K : \mathbb{R}^N \times \mathbb{R}^{\mathcal{D}} \rightarrow \mathbb{R}$, i.e.,

$$\tilde{y}_j(x, b) = \sum_{k=1}^K a_j^k \eta_j^k(x, b), \quad j = 1, \dots, J. \quad (24)$$

We consider two different ways to compute the \mathcal{F}_0 -measurable coefficients a_j^k , a variant of least-squares Monte Carlo, which picks up some ideas of the regression later approach of [13] in the optimal stopping literature, and a direct martingale minimization approach which builds on the works by [2] and [9] for optimal stopping.

As a key assumption on the basis functions, we impose that the expectations

$$R_j^k(x) := \left(\int_{\mathbb{R}^{\mathcal{D}}} b_1 \eta_{j+1}^k(x, b) P_{B_{j+1}}(db), \dots, \int_{\mathbb{R}^{\mathcal{D}}} b_D \eta_{j+1}^k(x, b) P_{B_{j+1}}(db) \right)^\top, \quad (25)$$

$x \in \mathbb{R}^N$, are available in closed form (or, can be computed numerically up to a ‘negligible’ error), cp. Remark 5.1 below. Defining $\tilde{Y}_j = \tilde{y}_j(X_{j-1}, B_j)$ as an approximation to Y_j , we, thus, observe that $E_j[\beta_{j+1} \tilde{Y}_{j+1}] = \sum_k a_{j+1}^k R_j^k(X_j)$ is given in closed form as well. Based on the input approximation \tilde{y} , we can now derive first approximations of the optimal control and of the Doob martingale M^* , from which input sub- and supersolutions can be obtained via the pathwise recursions (8) and (16). To compute such a control \tilde{r} , we solve (7) with Y_j replaced by \tilde{Y}_j , i.e. the process \tilde{r} is given by

$$\tilde{r}_j^\top E_j[\beta_{j+1} \tilde{Y}_{j+1}] - f_j^\#(\tilde{r}_j) = f_j(X_j, E_j[\beta_{j+1} \tilde{Y}_{j+1}]), \quad (26)$$

for every $j = 0, \dots, J - 1$, and belongs to \mathcal{A}_0 thanks to Lemma A.1 in [5]. Here, the convex conjugate can, of course, be approximated numerically as well. For the first approximation of the Doob martingale M^* , we, again, just replace the true solution Y by its approximation \tilde{Y} . Hence, a first approximation \tilde{M} of M^* is given by

$$\tilde{M}_j = \sum_{i=0}^{j-1} \beta_{i+1} \tilde{Y}_{i+1} - E_i[\beta_{i+1} \tilde{Y}_{i+1}], \quad j = 0, \dots, J. \quad (27)$$

Plugging \tilde{r} and \tilde{M} into the recursions (8) and (16) for θ^{low} and θ^{up} , we obtain the input sub- and supersolutions $Y_j^{(low,0)} = E_j[\theta_j^{low}(\tilde{r}, \tilde{M})]$ and $Y_j^{(up,0)} = E_j[\theta_j^{up}(\tilde{M})]$.

We can now sample Λ^{out} independent copies $(B_j(\lambda^{out}), j = 1, \dots, J)_{\lambda^{out}=1, \dots, \Lambda^{out}}$ of B , to which we refer as ‘outer’ paths. Then, we can compute the pathwise recursions for $\theta^{low}(\tilde{r}, \tilde{M})$ and $\theta^{up}(\tilde{M})$ along each of these outer paths and denote them by $\theta_j^{(low,0)}(\lambda^{out})$ and $\theta_j^{(up,0)}(\lambda^{out})$, $j = 0, \dots, J$, respectively. Applying the plain Monte Carlo estimator

$$\hat{Y}_0^{(up,0)} := \frac{1}{\Lambda^{out}} \sum_{\lambda^{out}=1}^{\Lambda^{out}} \theta_0^{(up,0)}(\lambda^{out}) \quad (28)$$

for $E_0[\theta_0^{up}(\tilde{M})]$ and the associated empirical standard deviation, one can compute an (asymptotic) confidence interval for $E_0[\theta_0^{up}(\tilde{M})]$ and thus an upper confidence bound on $Y_0 \leq E_0[\theta_0^{up}(\tilde{M})]$, see Section 1.1.3 of [12]. Analogously, from $(\theta_0^{(low,0)}(\lambda^{out}))_{\lambda^{out}=1, \dots, \Lambda^{out}}$, a lower confidence bound can be constructed, and, combining both bounds, we end up with an asymptotic confidence interval. We emphasize that Y_0 is a deterministic real number, but the construction of the confidence interval is conditional on any set of sample paths which might be used to pre-compute the coefficients a_j in (24) and which we think of as being included in \mathcal{F}_0 .

When such a confidence interval is not yet sufficiently tight for the application under consideration, one can run the iterative improvement algorithm described in Section 5.2 below. We shall first, however, discuss two ways to obtain the coefficients for the input approximation (24).

Least-squares Monte Carlo approach

The idea of least-squares Monte Carlo (LSMC) is to approximate the conditional expectation in (23) by an orthogonal projection onto a set of basis functions via regression, i.e., one computes

$$\tilde{Y}_j = f_j(X_j, \mathcal{P}_j[\beta_{j+1} \tilde{Y}_{j+1}]), \quad j = 0, \dots, J - 1, \quad \tilde{Y}_J = g(X_J),$$

as an approximation to Y , where \mathcal{P}_j denotes the empirical regression (given a set of sample paths) on a pre-specified basis. Note that one actually has to calculate D empirical regressions in each time step, since the stochastic weight β is \mathbb{R}^D -valued, and that the expression $\beta_{j+1} \tilde{Y}_{j+1}$ may suffer from a large variance, e.g., in the BSDE case (3), where the variance of the Malliavin Monte Carlo weights β for the first space derivative explodes as the time discretization becomes finer and finer. With our standing assumption (25) on the basis functions we can, instead, implement the following single-regression variant of least-squares Monte Carlo:

$$\tilde{Y}_j = \mathcal{P}_j \left[f_j(X_j, E_j[\beta_{j+1} \tilde{Y}_{j+1}]) \right], \quad j = 0, \dots, J - 1, \quad \tilde{Y}_J = \mathcal{P}_J[g(X_J)],$$

as (inductively) \tilde{Y}_{j+1} is a linear combination of $(\eta_{j+1}^k(X_j, B_{j+1}))_{k=1, \dots, K}$ and, thus, the conditional expectation inside f_j is available in closed form. This idea to ‘regress later’ originates in [13] for optimal stopping and was extended to the time discretization of BSDEs in [3], where a tremendous variance reduction effect is observed in the numerical examples.

To be more formal, we assume that Λ^{reg} independent copies of B to which we refer as ‘regression paths’ are given. The trajectories of β and of the Markovian process X along the λ th regression path are denoted by $\beta(\lambda)$ and $X(\lambda)$, $\lambda = 1, \dots, \Lambda^{reg}$. For the initialization of our algorithm, we require an approximation of the terminal condition $g_J(X_J) = y_J(X_{J-1}, B_J)$ in terms of basis functions. Applying a standard regression approach, we compute \mathbb{R}^K -valued coefficients $a_J = (a_J^1, \dots, a_J^K)$ via

$$a_J = \operatorname{argmin}_{a \in \mathbb{R}^K} \frac{1}{\Lambda^{reg}} \sum_{\lambda=1}^{\Lambda^{reg}} \left(g(X_J(\lambda)) - \sum_{k=1}^K a^k \eta_J^k(X_{J-1}(\lambda), B_J(\lambda)) \right)^2$$

and obtain $\tilde{y}_J(x, b) = \sum_{k=1}^K a_J^k \eta_J^k(x, b)$ as an approximation of $y_J(x, b) = g(h(x, b))$. Now, assume that an approximation $\tilde{y}_{j+1}(x, b)$ in terms of the basis functions has already been computed, i.e. $\tilde{y}_{j+1}(x, b) = \sum_{k=1}^K a_{j+1}^k \eta_{j+1}^k(x, b)$, with \mathbb{R}^K -valued coefficients $a_{j+1} = (a_{j+1}^1, \dots, a_{j+1}^K)$. Then, by (25),

$$f_j \left(X_j, E_j \left[\beta_{j+1} \sum_{k=1}^K a_{j+1}^k \eta_{j+1}^k(X_j, B_{j+1}) \right] \right) = f_j \left(X_j, \sum_{k=1}^K a_{j+1}^k R_j^k(X_j) \right),$$

(where we, of course, formally, perform an initial enlargement of the filtration by the regression paths, which are assumed to be independent of (X, β)). Projecting the right-hand side empirically on the basis functions $(\eta_j^1, \dots, \eta_j^K)$ leads to

$$a_j = \operatorname{argmin}_{a \in \mathbb{R}^K} \frac{1}{\Lambda^{reg}} \sum_{\lambda=1}^{\Lambda^{reg}} \left(f_j(X_j(\lambda), \sum_{k=1}^K a_{j+1}^k R_j^k(X_j(\lambda))) - \sum_{k=1}^K a^k \eta_j^k(X_{j-1}(\lambda), B_j(\lambda)) \right)^2,$$

and, thus, we obtain $\tilde{y}_j(x, b) = \sum_{k=1}^K a_j^k \eta_j^k(x, b)$ as an approximation to y_j .

Remark 5.1. In contrast to [13] and [3], we merely require in (25) that conditional expectations are available explicitly one step ahead, while [13] and [3] both additionally assume that the basis functions form martingales, i.e., $E_j[\eta_{j+1}^k(X_j, B_{j+1})] = \eta_j^k(X_{j-1}, B_j)$. One can exploit the additional flexibility in the following way: Suppose that each basis function η_j^k can be written in the product form $\eta_j^k(x, b) = \eta_j^{k,1}(x) \eta_j^{k,2}(h_j(x, b))$ and assume that $E[\beta_j \eta_j^{k,2}(h_j(x, B_j))]$ is available in closed form. Then, the expression in (25) is also available in closed form, as required. In particular, while the choice of $\eta_j^{k,2}$ is restricted to functions where explicit computations are possible, we are completely flexible in capturing a more complex dependence on the process X through the factor $\eta_j^{k,1}(x)$. In the numerical example of Section 6, we illustrate such a choice of basis functions.

Martingale minimization approach

It has been observed in the context of optimal stopping and in the BSDE examples in [4] and [5] that the construction of tight supersolutions can be significantly more difficult than the construction of tight subsolutions. The idea of the martingale minimization approach is thus to compute

the coefficients in (24) in such a way that the upper bound implied by the input supersolution is minimized. In the context of optimal stopping, similar ideas have been developed in [9] and [2]. In contrast to least-squares Monte Carlo, the optimization is now global and so the coefficients in (24) do not depend on the time index j .

In view of (25) each basis function defines a martingale via $M_0^{\{k\}} = 0$ and

$$M_j^{\{k\}} - M_{j-1}^{\{k\}} = \beta_j \eta_j^k(X_{j-1}, B_j) - R_{j-1}^k(X_{j-1}), \quad k = 1, \dots, K.$$

Writing,

$$M_j^a = \sum_{k=1}^K a^k M_j^{\{k\}}, \quad j = 0, \dots, J, \quad (29)$$

where $a = (a^1, \dots, a^K) \in \mathbb{R}^K$, we wish to choose a coefficient vector a^* , for which $E_0[\theta_0^{up}(M^a)]$ becomes minimal. Taking the pathwise optimality of the optimal martingale M^* in (18) into account and following the approach analyzed in [2] for optimal stopping, we add a standard deviation penalty to this minimization problem. To make the approach implementable, the expectation and standard deviation need to be replaced by empirical estimators over sample paths. To this end, we sample Λ^{mini} independent copies of B (which we refer to as ‘minimization paths’) and denote the evaluation of β and X along the λ th minimization path by $\beta(\lambda)$ and $X(\lambda)$. We then solve for

$$a^* = \operatorname{argmin}_{a \in \mathbb{R}^K} \left(\hat{E}_0[\theta_0^{up}(M^a)] + \gamma \sqrt{\frac{1}{\Lambda^{mini} - 1} \sum_{\lambda=1}^{\Lambda^{mini}} \left(\theta_0^{up}(M^a; \lambda) - \hat{E}[\theta_0^{up}(M^a)] \right)^2} \right), \quad (30)$$

where $\gamma \geq 0$ is fixed,

$$\hat{E}_0[\theta_0^{up}(M^a)] = \frac{1}{\Lambda^{mini}} \sum_{\lambda=1}^{\Lambda^{mini}} \theta_0^{up}(M^a; \lambda),$$

and $\theta_0^{up}(M^a; \lambda)$ is sampled according to (16) along the λ th minimization path. In our numerical examples, we use the MATLAB implementation of the Nelder-Mead simplex algorithm to search for a^* . An approximation to y_j as in (24) is then given by

$$\tilde{y}_j(x, b) = \sum_{k=1}^K a^{k,*} \eta_j^k(x, b), \quad j = 1, \dots, J.$$

Remark 5.2. The minimization approach requires the choice of the parameter γ . In our numerical results presented in Section 6, we apply a “training and testing” approach to tune this parameter. To this end, we choose a set $\{\gamma_1, \dots, \gamma_L\}$, $L \in \mathbb{N}$, of parameters. For each γ_l , $l = 1, \dots, L$, we compute a vector of coefficients $a_{\gamma_l}^* \in \mathbb{R}^K$ according to (30) along the minimization paths Λ^{mini} . If vectors $a_{\gamma_1}^*, \dots, a_{\gamma_L}^*$ are computed, we sample a new set of Λ^{test} test paths (independent copies of B which are also independent of the minimization paths). The parameter γ is obtained by taking the γ_l such that $a_{\gamma_l}^*$ minimizes the expression in brackets on the right hand side of (30) along the test paths over the set $\{a_{\gamma_1}^*, \dots, a_{\gamma_L}^*\}$. We note that in our experience the method’s practical performance is not particularly sensitive to the choice of γ .

5.2 Iterative improvement algorithm

We now assume that we are given input super- and subsolutions of the form $Y_j^{(up,0)} = E_j[\theta^{up}(\tilde{M})]$ and $Y_j^{(low,0)} = E_j[\theta^{low}(\tilde{r}, \tilde{M})]$ such that the control $\tilde{r} \in \mathcal{A}_0$ and the martingale $\tilde{M} \in \mathcal{M}_D$ can be evaluated in closed form along a given path B , cp. the constructions in Section 5.1. In order to compute the first iteration $\theta^{(low,1)}$ in (12) and $\theta^{(up,1)}$ in (20) we need to approximate the conditional expectations $E_j[\beta_{j+1}\theta_{j+1}^{low}(\tilde{r}, \tilde{M})]$, $E_j[\beta_{j+1}\theta_{j+1}^{up}(\tilde{M})]$, and $E_{j+1}[\theta_{j+1}^{up}(\tilde{M})]$. In the following, we focus on the supersolution case, but note that the subsolution case is analogous.

In our plain Monte Carlo implementation, we first sample Λ^{out} independent copies $B(\lambda^{out})$, $\lambda^{out} = 1, \dots, \Lambda^{out}$, of B . Moreover, for every time step j and outer path $B(\lambda^{out})$, we generate a new sample of independent copies $(B_i(\lambda^{mid}, j))_{i \geq j+1}$, $\lambda^{mid} = 1, \dots, \Lambda^{mid}$, of $(B_i)_{i \geq j+1}$. We denote by $B(\lambda^{out}, \lambda^{mid}, j)$ the path given by $(B_1(\lambda^{out}), \dots, B_j(\lambda^{out}), B_{j+1}(\lambda^{mid}, j), \dots, B_J(\lambda^{mid}, j))$, which switches from a given outer path to the corresponding middle path at time $j+1$. Similarly to the notation introduced before, we write $\beta(\lambda^{out}, \lambda^{mid}, j)$ and $\theta^{(up,0)}(\lambda^{out}, \lambda^{mid}, j)$ for the trajectories of β and $\theta^{up}(\tilde{M})$ along the path $B(\lambda^{out}, \lambda^{mid}, j)$. Along each outer path, we approximate the martingale $M^{(1)}$ in (21) with increment

$$M_{j+1}^{(1)} - M_j^{(1)} = \beta_{j+1}E_{j+1}[\theta_{j+1}^{(up,0)}] - E_j[\beta_{j+1}\theta_{j+1}^{(up,0)}], \quad j = 0, \dots, J-1,$$

by the plain Monte Carlo estimator

$$\tilde{M}_{j+1}^{(1)}(\lambda^{out}) - \tilde{M}_j^{(1)}(\lambda^{out}) = \beta_{j+1}(\lambda^{out})\hat{E}_{j+1}[\theta_{j+1}^{(up,0)}](\lambda^{out}) - \hat{E}_j[\beta_{j+1}\theta_{j+1}^{(up,0)}](\lambda^{out})$$

where

$$\begin{aligned} \hat{E}_j[\theta_j^{(up,0)}](\lambda^{out}) &:= \frac{1}{\Lambda^{mid}} \sum_{\lambda^{mid}=1}^{\Lambda^{mid}} \theta_j^{(up,0)}(\lambda^{out}, \lambda^{mid}, j) \\ \hat{E}_j[\beta_{j+1}\theta_{j+1}^{(up,0)}](\lambda^{out}) &:= \frac{1}{\Lambda^{mid}} \sum_{\lambda^{mid}=1}^{\Lambda^{mid}} \beta_{j+1}(\lambda^{out}, \lambda^{mid}, j)\theta_{j+1}^{(up,0)}(\lambda^{out}, \lambda^{mid}, j). \end{aligned} \quad (31)$$

Standard calculations show that $\tilde{M}^{(1)}$ is also a martingale when the filtration is suitably enlarged by the middle paths. We now write $\theta^{(up,1)}(\lambda^{out})$ for the realization of $\theta^{up}(\tilde{M}^{(1)})$ along the λ^{out} th outer path. Proceeding as in (28) ff., we can compute a new upper confidence bound for Y_0 based on $(\theta^{(up,1)}(\lambda^{out}))_{\lambda^{out}=1, \dots, \Lambda^{out}}$. Since $\tilde{M}^{(1)}$ converges to $M^{(1)}$ (along each outer path) as the number of middle paths converges to infinity, and since $E_0[\theta^{up}(M^{(1)})] \leq E_0[\theta^{up}(\tilde{M})]$ by Theorem 4.1, the corresponding upper bound is typically tighter than the one constructed from $(\theta^{(up,0)}(\lambda^{out}))_{\lambda^{out}=1, \dots, \Lambda^{out}}$, when the number of middle paths is sufficiently large.

If one wishes to compute a second iteration step (e.g., because the once improved confidence interval is still not tight enough), one can repeat this procedure with the only difference that we cannot assume the input martingale (which now is $\tilde{M}^{(1)}$) to be available in closed form along a given path. Its evaluation actually requires one layer of nested simulation as described above. However, in the next iteration step $\tilde{M}^{(1)}$ must be evaluated along middle paths and not along outer paths, and so the sampling of a third layer of Λ^{in} ‘inner paths’ is required. We do not get into any more details of the straightforward implementation, but note that, analogously, a third layer of simulation must already be sampled in the first iteration step, when the input martingale \tilde{M} is not available in closed form (e.g., when we drop assumption (25)).

In order to reduce the number of middle paths (in the first iteration step) and inner paths (in the second iteration step), we suggest to apply control variates in the plain Monte Carlo estimation (31) of the martingale increments based on the closed form expression for $E_j[\theta_{j+1}^{up}(\tilde{M})]$ and $E_j[\beta_{j+1}\theta_{j+1}^{up}(\tilde{M})]$. In our actual implementation, we proceed as described at the end of Section 4.1 in [4].

In principle, the algorithm can be further iterated, but each iteration step adds an additional layer of simulations. So, for practical reasons, we recommend not to run the algorithm with more than three layers of simulations, but rather to put more effort into the construction of the input approximations, when the confidence interval is still not tight enough. In our numerical test example below, very satisfactory 95% confidence bounds can be obtained with two iteration steps, even when the input approximation \tilde{y} in (24) is pre-computed by the martingale minimization approach with a single constant basis function.

6 Numerical example

In this section, we apply our approach to the problem of pricing a European option under funding constraints, i.e., under different interest rates for borrowing and lending. In the finance literature, this problem goes back to [6]. [19] emphasizes the relevance of such models in the light of the recent financial crisis. The model is also prominent example in the literature on backward stochastic differential equations starting with [10] and a well-established numerical test case [14, 20, 3, 4]. We begin by setting up the problem and explaining how it fits into our framework. Then, we present our numerical results. For the computation of input approximations we present different approaches, which incorporate a priori knowledge about the problem to a varying extent. The upper and lower bounds as well as the corresponding improvements are computed relying on the methodology of Section 5, including the use of control variates.

Pricing under funding risk

Let $0 = t_0 < t_1 < \dots < t_J = T$ be an equidistant discretization of the interval $[0, T]$. There are two riskless interest rates $R^l < R^b \in \mathbb{R}$ for lending respectively borrowing and N risky assets given by geometric Brownian motions X_1, \dots, X_N with dynamics

$$X_{n,j} = x_{n,0} \exp \left\{ \left(\mu - \frac{1}{2} \sum_{l=1}^N \sigma_{n,l}^2 \right) t_j + \sum_{l=1}^N \sigma_{n,l} W_{l,t_j} \right\}, \quad n = 1, \dots, N,$$

at t_j , for $j = 0, \dots, J$. Here, $x_{n,0}, \mu \in \mathbb{R}$, σ is an invertible $N \times N$ -matrix with entries in \mathbb{R} and W_1, \dots, W_N are independent Brownian motions. We consider the problem of pricing a European option on the assets X_1, \dots, X_N with maturity T and payoff $g(X_{1,J}, \dots, X_{N,J})$. When g satisfies a polynomial growth condition, the option payoff belongs to $L_J^{\infty-}(\mathbb{R})$. Applying the discretization scheme proposed in [11] to equation (1.11) in [10], the value Y of the option on the time grid $\{t_0, \dots, t_J\}$ is given by

$$Y_j = (1 - R^l \Delta) E_j[Y_{j+1}] - (\mu - R^l) Z_j^\top \sigma^{-1} \mathbf{1} \Delta + (R^b - R^l) \Delta (E_j[Y_{j+1}] - Z_j^\top \sigma^{-1} \mathbf{1})_-, \quad (32)$$

with terminal condition $Y_J = g(X_{1,J}, \dots, X_{N,J})$. Here, $\Delta := t_j - t_{j-1}$ for $j = 1, \dots, J$, $\mathbf{1} \in \mathbb{R}^N$ is a vector consisting of ones, and $(x)_- := \max\{-x, 0\}$ for $x \in \mathbb{R}$. Moreover, the random vector Z_j is

given by

$$Z_{n,j} := E_j \left[\frac{\Delta W_{n,j+1}}{\Delta} Y_{j+1} \right], \quad n = 1, \dots, N,$$

where $E_j[\cdot]$ denotes the conditional expectation with respect to the information generated by the multivariate Brownian motion up to time t_j , and $\Delta W_{n,j+1} := W_{n,t_{j+1}} - W_{n,t_j}$. The term $Z_j^\top \sigma^{-1} \mathbf{1}$ in (32) represents the overall position in the risky assets in the hedging portfolio at time t_j . Therefore, $E_j[Y_{j+1}] - Z_j^\top \sigma^{-1} \mathbf{1}$ is an approximation of the position in the bank account at time t_j . The sign of this expression determines which interest rate is applicable. By taking the function $F_j : \mathbb{R}^D \rightarrow \mathbb{R}$, $D = N + 1$, given by

$$F_j(z) = (1 - R^l \Delta) z_0 - (\mu - R^l) z_{(-0)}^\top \sigma^{-1} \mathbf{1} \Delta + (R^b - R^l) \Delta \left(z_0 - z_{(-0)}^\top \sigma^{-1} \mathbf{1} \right)_-,$$

where $z_{(-0)} := (z_1, \dots, z_N)$, and setting

$$B_{j+1} = \beta_{j+1} = \left(1, \frac{p_C(\Delta W_{1,j+1})}{\Delta}, \dots, \frac{p_C(\Delta W_{N,j+1})}{\Delta} \right)^\top, \quad j = 0, \dots, J-1,$$

we observe that the recursion (32) fits into our framework. Here, p_C denotes a truncation function, i.e. $p_C(x) = -C \vee x \wedge C$ for $C \in \mathbb{R}_+$. Notice that the truncation of $W_{t_{j+1}} - W_{t_j} \sim \mathcal{N}(0, \Delta)$ at C becomes arbitrarily mild as Δ gets small. Truncation is required, because the increments of the Brownian motions are unbounded and, thus, the monotonicity assumption (4) might be violated. A sufficient condition for (4) to hold is then

$$\Delta \cdot \max\{|R^l|, |R^b|\} + C \cdot \max\{|\mu - R^l|, |\mu - R^b|\} \cdot \sum_{d=1}^D \left| \sum_{l=1}^D (\sigma^{-1})_{d,l} \right| \leq 1, \quad (33)$$

see [20] for an analysis of this type of truncation error.

Recall that our algorithm requires the convex conjugate $F_j^\#$ of F_j when running the pathwise recursion formula for θ^{low} and it also requires to solve for the optimality condition (13) in the iteration for the lower bounds. As F is piecewise linear, the convex conjugate is straightforward to compute and equals $F_j^\# \equiv 0$ on its effective domain. Moreover, with the function $u : \mathbb{R} \rightarrow \mathbb{R}^{N+1}$ given by

$$u^{(0)}(s) = (1 - s\Delta) \quad \text{and} \quad u^{(n)}(s) = -(\mu - s)\Delta \sum_{l=1}^N (\sigma^{-1})_{n,l}, \quad n = 1, \dots, N,$$

the effective domain of the convex conjugate is $D_{F^\#}^{(j,\omega)} = \{u(R) | R \in [R^l, R^b]\}$. Finally, for every $z = (z_0, \dots, z_N) \in \mathbb{R}^{N+1}$,

$$r(z) = \begin{cases} u(R^l), & z_0 \geq (z_1, \dots, z_N) \sigma^{-1} \mathbf{1} \\ u(R^b), & z_0 < (z_1, \dots, z_N) \sigma^{-1} \mathbf{1} \end{cases}$$

solves $r(z)^\top z = F(z)$, compare (13).

Benchmark product

For our numerical experiments, we consider the example discussed in [4], but add a non-trivial correlation structure to the problem. This example is a multidimensional version of an example going back to [14]. We compute upper and lower bounds on the price of a European call-spread option with strikes K_1 and K_2 on the maximum of five assets, i.e.,

$$g(x_1, \dots, x_5) = \left(\max_{d=1, \dots, 5} x_d - K_1 \right)_+ - 2 \left(\max_{d=1, \dots, 5} x_d - K_2 \right)_+, \quad x \in \mathbb{R}^5.$$

The maturity T is set to three months, i.e. $T = 0.25$, and the strikes are $K_1 = 95$ and $K_2 = 115$. The interest rates R^l and R^b are 1% and 6%. For the geometric Brownian motions X_1, \dots, X_5 we take $x_{d,0} = 100$, $d = 1, \dots, 5$, as starting value and choose the drift μ to be 0.05. In contrast to [4] we do not assume that X_1, \dots, X_5 are independent and consider the diffusion matrix σ given by

$$\sigma = \tilde{\sigma} \cdot \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ \rho & \sqrt{1-\rho^2} & 0 & 0 & 0 \\ \rho & 0 & \sqrt{1-\rho^2} & 0 & 0 \\ \rho & 0 & 0 & \sqrt{1-\rho^2} & 0 \\ \rho & 0 & 0 & 0 & \sqrt{1-\rho^2} \end{pmatrix},$$

where $\tilde{\sigma} = 0.2$. In our numerical experiments below, the correlation parameter ρ varies in the interval $[-0.3, 0.3]$ and the time discretization J takes values in $\{20, 30, 40\}$. With this choice of parameters, we observe that (33) holds with $C = 0.77$ at the roughest time discretization level $J = 20$. Truncating the Brownian increments with standard deviation $\sqrt{\Delta} \approx 0.112$ at 0.77 is the same as truncating a standard normal random variable at 6.88, corresponding to truncating a probability mass of $3 \cdot 10^{-12}$ in both tails.

Generic minimization algorithm

For the construction of the input approximation, we first run the martingale minimization algorithm with the single and completely generic basis function $\eta_j^1(x, b) := 1$, i.e., we initially approximate Y_j by a constant and the $Z_{n,j}$ by zero, $n = 1, \dots, N$. Then, in the minimization approach presented in Section 5.1 we have a single 6-dimensional martingale $M^{\{1\}}$ given by $\tilde{M}_{0,j+1}^{\{1\}} - \tilde{M}_{0,j}^{\{1\}} = 0$ and

$$\tilde{M}_{d,j+1}^{\{1\}} - \tilde{M}_{d,j}^{\{1\}} = \beta_{d,j+1} - E_j[\beta_{d,j+1}] = \frac{p_C(\Delta W_{d,j+1})}{\Delta}$$

for $d = 1, \dots, 5$. In order to compute the \mathbb{R} -valued coefficient a^* , and, hence, the constant approximation $\tilde{y}_j(x, b) = a^*$ to y_j , we implement the ‘training and testing’ approach of Remark 5.2 with $\Lambda^{mini} = \Lambda^{test} = 1000$ paths and $\{\gamma_1, \dots, \gamma_{21}\} = \{0, 0.025, \dots, 0.5\}$. We find that a^* , as an approximation of Y_0 , ranges between 16 and 17.5 for our different choices of J and ρ , and as $a^* > 0$, the input subsolution $Y^{(low,0)}$ is constructed from the constant control $u(R^l)$. For the computation of upper and lower bounds with up to two iterative improvements, we take $\Lambda^{out} = 1000$ outer paths, $\Lambda^{mid} = 200$ middle paths and $\Lambda^{in} = 50$ inner paths. The resulting estimators for the upper and lower bounds from the k -th improvement are denoted by $\hat{Y}_0^{(up,k),a}$ and $\hat{Y}_0^{(low,k),a}$. For comparison, we also state the upper bound estimator $\hat{Y}_0^{(up,0),0}$ which is computed by choosing $a = 0$, i.e., by setting all martingale increments to zero.

ρ	0.3			-0.3		
J	20	30	40	20	30	40
$\hat{Y}_0^{(up,0),0}$	18.9637 (0.2243)	20.5682 (0.2444)	22.1942 (0.2720)	26.1759 (0.2736)	29.5829 (0.3217)	34.2500 (0.3424)
$\hat{Y}_0^{(up,0),a^*}$	14.4278 (0.1405)	14.5533 (0.1330)	14.7838 (0.1343)	15.9557 (0.1081)	16.0660 (0.1025)	16.5631 (0.0948)
$\hat{Y}_0^{(up,1),a^*}$	13.1430 (0.0129)	13.2626 (0.0133)	13.3451 (0.0154)	14.5063 (0.0127)	14.6878 (0.0148)	14.9476 (0.0136)
$\hat{Y}_0^{(up,2),a^*}$	13.0461 (0.0127)	13.1088 (0.0137)	13.0919 (0.0139)	14.2047 (0.0107)	14.2452 (0.0108)	14.3340 (0.0102)
$\hat{Y}_0^{(low,0),a^*}$	12.6157 (0.0231)	12.6281 (0.0253)	12.5792 (0.0307)	13.7919 (0.0289)	13.7291 (0.0366)	13.8283 (0.0368)
$\hat{Y}_0^{(low,1),a^*}$	12.9915 (0.0139)	13.0063 (0.0150)	12.9703 (0.0184)	14.0492 (0.0180)	14.0000 (0.0227)	14.0498 (0.0233)

Table 1: Upper and lower bounds based on the generic minimization algorithm for different time discretizations. Standard deviations are given in brackets.

Table 1 presents upper and lower bounds for two different choices of ρ , namely $\rho = 0.3$ and $\rho = -0.3$.

We first observe that the upper bound is very sensitive with respect to the input martingale. Even optimizing a very crude constant approximation for Y has a huge impact, and, e.g., leads to a half as large upper bound for $J = 40$ time steps in the negative correlation case compared to the upper bound computed from the zero martingale $\hat{Y}_0^{(up,0),0}$. Nonetheless, the relative width of the 95% confidence interval based on the optimal constant approximation is still more than 20% for 40 time steps in the positive correlation case and even larger in the negative correlation case. Improving upper and lower confidence bound once, shrinks the 95% confidence interval to a quite acceptable relative width of less than 3.5% in the positive correlation case, while a second iterative improvement of the upper bound leads to relative width of less than 1.5%. The negative correlation apparently makes the problem harder to solve numerically. But, still, after two iteration steps for the upper bound and one iteration step for the lower bound we end up with a 95% confidence interval of a relative width of less than 2.5%. We also observe a significant decrease in the empirical standard deviations of the upper bound estimators through the improvement steps, as expected since the martingales approach the pathwise optimal Doob martingale of βY , cp. (18).

Taking into account that no problem-specific information was used to construct the above confidence intervals in a five-dimensional problem with non-smooth coefficients and non-trivial correlation structure, we believe that the numerical results are rather striking. We note, however, that the second iteration step increases the computational costs by a factor of $\Lambda^{in} \cdot (J/3)$ (e.g., a factor of 667 in our setting for $J = 40$ time steps) compared to a single improvement step. Thus, we next explore to what extent the results can be improved by putting more effort into the construction of the input approximation.

Non-generic minimization and LSMC algorithms

Following ideas of [1] for the pricing of Bermudan options on the maximum of several assets, we now incorporate information about option prices on the largest and second-largest asset into the

function basis. To this end, we define the two adapted processes $d^{(1)}$ and $d^{(2)}$ by

$$\begin{aligned} d_j^{(1)} &:= \inf \{d \in \{1, \dots, 5\} \mid X_{d,j} \geq X_{n,j} \ \forall n = 1, \dots, 5\} \\ d_j^{(2)} &:= \inf \left\{ d \in \{1, \dots, 5\} \setminus \{d_j^{(1)}\} \mid X_{d,j} \geq X_{n,j} \ \forall n \in \{1, \dots, 5\} \setminus \{d_j^{(1)}\} \right\} \end{aligned}$$

for $j = 0, \dots, J$. Hence, $d_j^{(1)}$ and $d_j^{(2)}$ indicate the largest respectively second-largest asset at time t_j . In particular, they can be viewed as functions of X_j . Based on this, we define the following functions which serve as a basis for our approximations of Y :

$$\begin{aligned} \eta_j^1(X_{j-1}, X_j) &:= 1, \quad \eta_j^{\iota+1}(X_{j-1}, X_j) := \sum_{d=1}^5 X_{d,j} 1_{\{d_{j-1}^{(\iota)}=d\}}, \quad \iota = 1, 2, \\ \eta_j^{\iota+3}(X_{j-1}, X_j) &:= \sum_{d=1}^5 E \left[(X_{d,J} - K_1)_+ - 2(X_{d,J} - K_2)_+ \mid X_{d,j} \right] 1_{\{d_{j-1}^{(\iota)}=d\}}, \quad \iota = 1, 2, \\ \eta_j^6(X_{j-1}, X_j) &:= \sum_{d=1}^5 E \left[(X_{d,J} - K_2)_+ \mid X_{d,j} \right] 1_{\{d_{j-1}^{(1)}=d\}}. \end{aligned}$$

For $j = 0$, we replace $1_{\{d_{j-1}^{(\iota)}=d\}}$ by $1_{\{d_j^{(\iota)}=d\}}$, $\iota = 1, 2$. Here, we write, for simplicity and in slight abuse of notation, the basis functions as functions of (X_{j-1}, X_j) instead of (X_{j-1}, B_j) . Note that, e.g., the fourth basis function represents the price of the corresponding call spread option at time t_j on the asset which is the largest one at time t_{j-1} . Shifting the time index in the indicator by one time step (compared to the more intuitive function basis in [1] which is based on the largest asset at time t_j) turned out to be inessential in this numerical example, but ensures that the 'one-step' conditional expectations $R_{j-1}^k(X_{j-1})$ in (25) are available in closed form. These are, essentially, Black-Scholes prices and Black-Scholes deltas of European options at time t_{j-1} on the asset which is the (second) largest at time t_{j-1} .

With these basis functions, we construct input approximations as described in Section 5.1. For the martingale minimization algorithm, we run as before $\Lambda^{mini} = \Lambda^{test} = 1000$ paths and take the penalization parameter from the set $\{\gamma_1, \dots, \gamma_{21}\} = \{0, 0.025, \dots, 0.5\}$. The modified LSMC approach is applied with $\Lambda^{reg} = 100.000$ regression paths. Tables 2 and 3 below display the corresponding upper and lower bound estimators as well as iterative improvements up to the second order, based on these two input approximations. As before, we denote by $\hat{Y}_0^{(up,k)}$ and $\hat{Y}_0^{(low,k)}$ the upper respectively lower bound resulting from the k -th improvement.

By and large, we find that the quality of the upper bound estimators $\hat{Y}_0^{(up,0),mini}$ and $\hat{Y}_0^{(up,0),reg}$, computed from the two different methods to obtain the coefficients for the input approximation, is almost identical. They typically vary by less than two empirical standard deviations. The same holds true for the lower bounds $\hat{Y}_0^{(low,0),mini}$ and $\hat{Y}_0^{(low,0),reg}$. We also observe that, compared to the generic implementation, the input lower bounds $\hat{Y}_0^{(low,0),mini}$ and $\hat{Y}_0^{(low,0),reg}$ are of the same quality as the generic lower bounds in Table 1 $\hat{Y}_0^{(low,1),a^*}$ after one iterative improvement. Similarly, one improvement step of the upper bound in both non-generic cases $\hat{Y}_0^{(up,1),mini}$ and $\hat{Y}_0^{(up,1),reg}$ is comparable with two improvement steps in the generic setting $\hat{Y}_0^{(up,2),a^*}$. Recalling the large computational costs for the second improvement step, we observe that incorporating soft problem information into the function basis (here, the indicator function on the largest and second-largest asset one time step before) can significantly help to pin down the nonlinear option price

ρ	0.3			-0.3		
J	20	30	40	20	30	40
$\hat{Y}_0^{(up,0),mini}$	13.3465 (0.0694)	13.3766 (0.0738)	13.5420 (0.0741)	14.7198 (0.0658)	14.9629 (0.0676)	15.0104 (0.0634)
$\hat{Y}_0^{(up,1),mini}$	13.0424 (0.0065)	13.0595 (0.0072)	13.0738 (0.0076)	14.2064 (0.0065)	14.2828 (0.0064)	14.3567 (0.0071)
$\hat{Y}_0^{(up,2),mini}$	13.0423 (0.0070)	13.0768 (0.0070)	13.0751 (0.0071)	14.1761 (0.0060)	14.1939 (0.0056)	14.2293 (0.0057)
$\hat{Y}_0^{(low,0),mini}$	12.9953 (0.0076)	12.9737 (0.0102)	12.9923 (0.0098)	14.0466 (0.0112)	14.0773 (0.0093)	14.0754 (0.0123)
$\hat{Y}_0^{(low,1),mini}$	13.0167 (0.0068)	13.0171 (0.0076)	13.0101 (0.0082)	14.0835 (0.0075)	14.1011 (0.0075)	14.1035 (0.0093)

Table 2: Upper and lower bounds based on the non-generic minimization algorithm for different time discretizations. Standard deviations are given in brackets.

ρ	0.3			-0.3		
J	20	30	40	20	30	40
$\hat{Y}_0^{(up,0),reg}$	13.2481 (0.0654)	13.3234 (0.0660)	13.3730 (0.0694)	14.7348 (0.0720)	14.9905 (0.0726)	14.9565 (0.0782)
$\hat{Y}_0^{(up,1),reg}$	13.0439 (0.0061)	13.0479 (0.0057)	13.0675 (0.0056)	14.2022 (0.0060)	14.2704 (0.0063)	14.3315 (0.0064)
$\hat{Y}_0^{(up,2),reg}$	13.0503 (0.0064)	13.0681 (0.0065)	13.0857 (0.0067)	14.1840 (0.0057)	14.2207 (0.0056)	14.2351 (0.0059)
$\hat{Y}_0^{(low,0),reg}$	12.9958 (0.0070)	13.0059 (0.0079)	12.9979 (0.0086)	14.0320 (0.0113)	14.0584 (0.0099)	14.0638 (0.0129)
$\hat{Y}_0^{(low,1),reg}$	13.0171 (0.0068)	13.0192 (0.0075)	13.0118 (0.0082)	14.0864 (0.0074)	14.1046 (0.0077)	14.1090 (0.0093)

Table 3: Upper and lower bounds based on the modified LSMC algorithm for different time discretizations. Standard deviations are given in brackets.

Y_0 into a rather tight confidence interval after one iteration step only (and, hence, at moderate costs). For the sake of completeness, we also report the numerical results after performing a second iteration step for the upper bounds in the non-generic case. While in the case of negative correlation, we obtain a further improvement and end up with a confidence interval of a relative width of less than 1.5 % for $J = 40$ time steps, the situation for the positive correlation case is different. Here, the theoretical improvement of the upper bound is offset by the additional upward bias due to the small number of inner paths. In this case, however, the relative width of the 95% confidence interval is about 0.75% already after one iteration step, and, thus, any further improvement seems to be unnecessary for the option pricing problem under consideration.

References

- [1] L. Andersen and M. Broadie. Primal-dual simulation algorithm for pricing multidimensional American options. *Management Science*, 50(9):1222–1234, 2004.
- [2] D. Belomestny. Solving optimal stopping problems via empirical dual optimization. *Annals of Applied Probability*, 23(5):1988–2019, 2013.
- [3] C. Bender and J. Steiner. Least-squares Monte Carlo for backward SDEs. In R. Carmona, P. Del Moral, P. Hu, and N. Oudjane, editors, *Numerical Methods in Finance*, pages 257–289. Springer, 2012.
- [4] C. Bender, N. Schweizer, and J. Zhuo. A primal-dual algorithm for BSDEs. *Mathematical Finance*, Early View, DOI: 10.1111/mafi.12100, 2015.
- [5] C. Bender, C. Gärtner, and N. Schweizer. Pathwise dynamic programming. *Working paper available at www.math.uni-sb.de/ag-bender/bender_gaertner_schweizer.pdf*, 2016.
- [6] Y. Z. Bergman. Option pricing with differential interest rates. *Review of Financial Studies*, 8(2):475–500, 1995.
- [7] D. B. Brown, J. E. Smith, and P. Sun. Information relaxations and duality in stochastic dynamic programs. *Operations Research*, 58(4):785–801, 2010.
- [8] N. Chen and P. Glasserman. Additive and multiplicative duals for American option pricing. *Finance and Stochastics*, 11(2):153–179, 2007.
- [9] V. V. Desai, V. F. Farias, and C. C. Moallemi. Pathwise optimization for optimal stopping problems. *Management Science*, 58(12):2292–2308, 2012.
- [10] N. El Karoui, S. Peng, and M. C. Quenez. Backward stochastic differential equations in finance. *Mathematical Finance*, 7(1):1–71, 1997.
- [11] A. Fahim, N. Touzi, and X. Warin. A probabilistic numerical method for fully nonlinear parabolic PDEs. *Annals of Applied Probability*, 21(4):1322–1364, 2011.
- [12] P. Glasserman. *Monte Carlo Methods in Financial Engineering*. Springer, 2004.
- [13] P. Glasserman and B. Yu. Simulation for American options: Regression now or regression later? In H. Niederreiter, editor, *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pages 213–226. Springer, 2004.

- [14] E. Gobet, J.-P. Lemor, and X. Warin. A regression-based Monte Carlo method to solve backward stochastic differential equations. *Annals of Applied Probability*, 15(3):2172–2202, 2005.
- [15] E. Gobet, J. López-Salas, P. Turkedjiev, and C. Vázquez. Stratified regression Monte-Carlo scheme for semilinear PDEs and BSDEs with large scale parallelization on GPUs. *HAL Preprint hal-01186000*, 2015.
- [16] M. B. Haugh and L. Kogan. Pricing American options: A duality approach. *Operations Research*, 52(2):258–270, 2004.
- [17] R. A. Howard. *Dynamic Programming and Markov Processes*. Technology Press and Wiley, 1960.
- [18] A. Kolodko and J. Schoenmakers. Iterative construction of the optimal Bermudan stopping time. *Finance and Stochastics*, 10(1):27–49, 2006.
- [19] J.-P. Laurent, P. Amzelek, and J. Bonnaud. An overview of the valuation of collateralized derivative contracts. *Review of Derivatives Research*, 17(3):261–286, 2014.
- [20] J.-P. Lemor, E. Gobet, and X. Warin. Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. *Bernoulli*, 12(5):889–916, 2006.
- [21] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001.
- [22] E. Pardoux and S. Peng. Backward stochastic differential equations and quasilinear parabolic partial differential equations. In B. L. Rozovskii and R. B. Sowers, editors, *Stochastic partial differential equations and their applications*, pages 200–217. Springer, 1992.
- [23] M. L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 1994.
- [24] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- [25] L. C. G. Rogers. Monte Carlo valuation of American options. *Mathematical Finance*, 12(3):271–286, 2002.
- [26] J. N. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, 2001.